



# UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE  
United States Patent and Trademark Office  
Address: COMMISSIONER FOR PATENTS  
P.O. Box 1450  
Alexandria, Virginia 22313-1450  
[www.uspto.gov](http://www.uspto.gov)

lh

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/643,758	08/18/2003	Kitrick Sheets	1376.729US1	3924
21186	7590	03/10/2006	EXAMINER	
SCHWEGMAN, LUNDBERG, WOESSNER & KLUTH			TSAI, SHENG JEN	
1600 TCF TOWER			ART UNIT	PAPER NUMBER
121 SOUTH EIGHT STREET				2186
MINNEAPOLIS, MN 55402				

DATE MAILED: 03/10/2006

Please find below and/or attached an Office communication concerning this application or proceeding.

<b>Office Action Summary</b>	<b>Application No.</b>	<b>Applicant(s)</b>
	10/643,758 Sheng-Jen Tsai	SHEETS ET AL. Art Unit 2186

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --  
**Period for Reply**

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

#### Status

1) Responsive to communication(s) filed on 02 February 2006.  
 2a) This action is **FINAL**.                    2b) This action is non-final.  
 3) Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

#### Disposition of Claims

4) Claim(s) 1-17 is/are pending in the application.  
 4a) Of the above claim(s) \_\_\_\_\_ is/are withdrawn from consideration.  
 5) Claim(s) \_\_\_\_\_ is/are allowed.  
 6) Claim(s) 1-17 is/are rejected.  
 7) Claim(s) \_\_\_\_\_ is/are objected to.  
 8) Claim(s) \_\_\_\_\_ are subject to restriction and/or election requirement.

#### Application Papers

9) The specification is objected to by the Examiner.  
 10) The drawing(s) filed on \_\_\_\_\_ is/are: a) accepted or b) objected to by the Examiner.  
     Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).  
     Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).  
 11) The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

#### Priority under 35 U.S.C. § 119

12) Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).  
 a) All    b) Some \* c) None of:  
 1. Certified copies of the priority documents have been received.  
 2. Certified copies of the priority documents have been received in Application No. \_\_\_\_\_.  
 3. Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

\* See the attached detailed Office action for a list of the certified copies not received.

#### Attachment(s)

1) <input type="checkbox"/> Notice of References Cited (PTO-892)	4) <input type="checkbox"/> Interview Summary (PTO-413)
2) <input type="checkbox"/> Notice of Draftsperson's Patent Drawing Review (PTO-948)	Paper No(s)/Mail Date. _____ .
3) <input checked="" type="checkbox"/> Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08) Paper No(s)/Mail Date _____	5) <input type="checkbox"/> Notice of Informal Patent Application (PTO-152)
	6) <input type="checkbox"/> Other: _____

## DETAILED ACTION

1. This Office Action is taken in response to Applicants' Amendment and Remarks filed on February 2, 2006 regarding application 10,643,758 filed on August 18, 2003.

2. Claims 1, 4, 6, 9 and 11 have been amended.

Claims 12-17 have been added.

Claims 1-17 are pending in the application under consideration.

3. ***Response to Remarks and Amendments***

Applicants' amendments and remarks have been fully and carefully considered.

Independent claims 1, 4, 6, 9 and 11 have been amended to include the new limitation of "wherein exporting includes requesting, at a processor within each node, that the operating system load the RTT from the local address space of its respective node and requesting that the operating system enable remote translation."

In response to the amendments, a new ground of claim analysis, based on a reference (Schimmel, US 6,105,113) that was considered to be pertinent to Applicants' invention but not relied upon for claim analysis in the previous Office Action (cited in the ***Related Prior Art*** section), has been embarked. Refer to the corresponding sections of claim analysis for details.

### ***Claim Rejections - 35 USC § 102***

4. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

5. Claims 1-17 are rejected under 35 U.S.C. 102(b) as being anticipated by Schimmel (US 6,105,113).

As to claim 1, Schimmel discloses **a method of accessing shared memory in a computer system having a plurality of nodes** [System and Method for Maintaining Translation Look-aside Buffer (TLB) Consistency (title); figure 3 shows a Distributed Shared Memory (DSM) system], **including a first node** [any node in figure 3 may be the first node], **wherein each node includes a processor and local memory** [figure 3], **the method comprising:**

**distributing an application across the plurality of nodes** [the present invention can be implemented on any computer system that employs virtual memory. Thus, the present invention can be implemented in both uni-processor environments and multiple processor environments. The present invention is especially useful in shared memory, multi-processor systems where page migration occurs. Shared memory systems that benefit from the present invention include centralized shared memory systems, such as, symmetric multiple processor (SMP) systems and distributed shared memory (DSM) systems. The present invention can be employed to maintain consistency for any number of TLBs in a system (column 4, lines 32-43)];

**building an application virtual address space** [**in a virtual memory scheme**, each process that is allocated a block of physical memory is also provided with a set of translations for translating virtual addresses to assigned physical addresses of the

**allocated block ... (column 1, lines 24-46)], wherein building an application virtual address space includes:**

**building a local virtual address space for the application in each of the plurality of nodes, wherein the local virtual address space translates a virtual address generated by the application executing on that node to a physical address in local memory for that node [by distributing physical or main memory 328-342 throughout DSM 310, each processing node 350-364 can include a portion of main memory. This physical proximity between processor and memory reduces memory latency with respect to the processor and memory within a processing node. DSM 310 is preferably configured so that data which is accessed most frequently by a particular processing node is placed in the portion of main memory within the processing node. If that data is subsequently accessed more frequently by a processor in another processing node, the data is migrated, or moved, to a portion of main memory within the other processing node (column 7, lines 30-42); in operation, when a CPU requires a physical memory address that is associated with a virtual memory address, the CPU first searches the virtual address tag of the TLB table. If a valid translation is not found in the TLB table, the translation is retrieved from a cache or from main memory and a copy of the translation is placed in the TLB table (column 4, lines 8-13)]; and exporting the local virtual address space for each node to a Remote Translation Table (RTT) associated with that node [in a virtual memory scheme, each process that is allocated a block of physical memory is also provided with a set of translations for translating virtual addresses to assigned physical addresses of the allocated block.**

Each set of translations can be stored in, for example, a page table. A page table can be associated with a specific user or shared by multiple users (column 1, lines 24-31); figure 8 shows the RTT table; column 3, lines 42-62]; **wherein exporting includes requesting, at a processor within each node, that the operating system load the RTT from the local address space of its respective node and requesting that the operating system enable remote translation** [whenever a translation in TLB 718 is to be invalidated, CPU 714 or the operating system must be interrupted in order to execute the invalidation (column 9, lines 65-67); a copy of the translation is also placed in TLB 718. Later, when CPU 714 requires a translation, CPU 714 or an operating system searches TLB 718. If the translation is not found in TLB 718 (i.e., a TLB "miss"), the desired translation can be loaded from the page tables in memory by hardware, software, firmware, or any combination thereof (column 9, lines 22-28); an operating system that controls processor and cache node 410 can map virtual memory, as discussed above in FIGS. 5 and 6, where virtual memory addresses 512 are mapped to physical addresses 516 (column 11, lines 52-54); in step 912, the operating system generates virtual memory address-to-physical memory address translations for the mapped data. For example, a page table, such as page table 610 can be generated for a process that is provided with mapped memory. The page table can be for the exclusive use of one user or can be shared by multiple users (column 11, lines 59-64); the operation can be performed solely by the CPU, solely by an operating system (not shown) that controls the CPU or by a combination of the CPU and the operating system (column 10, lines 39-42); the PTE for page 0x10 can be found by

simply looking at index 0x10 in the page table array. The starting address of the array itself is maintained by the operating system in such a way that it is easy to find (column 8, lines 40-44)]; and

**performing a memory reference to a memory location in the application virtual address space, wherein performing a memory reference to a memory location in the application virtual address space includes translating bits of the application virtual address into a node address associated with the first node and translating bits of the application virtual address using the RTT associated with the first node** [in operation, when a CPU requires a physical memory address that is associated with a virtual memory address, the CPU first searches the virtual address tag of the TLB table. If a valid translation is not found in the TLB table, the translation is retrieved from a cache or from main memory and a copy of the translation is placed in the TLB table (column 4, lines 8-13); column 7, lines 50-67; column 8, lines 1-67].

As to claim 2, Schimmel teaches that **the local address space is read from a Translation Look-aside Buffer (TLB)** [figure 8; column 3, lines 42-62].

As to claim 3, In order for the remote translation mechanism disclosed by Schimmel to work and function properly, it is inherent that the RTT at all the nodes be initialized and synchronized first before any reference to a memory location resides at a remote node can be served. Without the initialization and synchronization, the RTT may not have the correct information to reach the correct memory location.

As to claim 4, Schimmel teaches a **system comprising:**  
**a plurality of nodes** [figure 3], **each node including:**

**one or more processors [figure 3];**  
**a memory [figure 3]; and**  
**a memory controller operatively coupled to the memory and the one or more processors** [figure 3 shows the cache coherency directory], **wherein the memory controller includes a Remote Translation Table (RTT)** [figure 8 shows the RTT table; column 3, lines 42-62], **wherein the RTT translates a virtual address received as part of a memory request received from another node into a memory request with physical addresses into the memory on the node associated with the RTT** [in operation, when a CPU requires a physical memory address that is associated with a virtual memory address, the CPU first searches the virtual address tag of the TLB table. If a valid translation is not found in the TLB table, the translation is retrieved from a cache or from main memory and a copy of the translation is placed in the TLB table (column 4, lines 8-13); column 7, lines 50-67; column 8, lines 1-67];  
**further wherein the RTT is initialized upon the start of a process associated with an application by building virtual to physical address translations for local virtual address space in the node corresponding to the application** [In order for the remote translation mechanism disclosed by Schimmel to work and function properly, it is inherent that the RTT at all the nodes be initialized and synchronized first before any reference to a memory location resides at a remote node can be served. Without the initialization and synchronization, the RTT may not have the correct information to reach the correct memory location], **and exporting the virtual to physical address translations for the local virtual address space from the node**

**to the Remote Translation Table (RTT) associated with that node** [in a virtual memory scheme, each process that is allocated a block of physical memory is also provided with a set of translations for translating virtual addresses to assigned physical addresses of the allocated block. Each set of translations can be stored in, for example, a page table. A page table can be associated with a specific user or shared by multiple users (column 1, lines 24-31); figure 8 shows the RTT table; column 3, lines 42-62].

As to claim 5, In order for the remote translation mechanism disclosed by Schimmel to work and function properly, it is inherent that the RTT at all the nodes be initialized and synchronized first before any reference to a memory location resides at a remote node can be served. Without the initialization and synchronization, the RTT may not have the correct information to reach the correct memory location.

As to claim 6, refer to "As to claim 1."

As to claim 7, refer to "As to claim 3."

As to claim 8, refer to "As to claim 2."

As to claim 9, refer to "As to claim 1" and "As to claim 4."

As to claim 10, refer to "As to claim 3."

As to claim 11, refer to "As to claim 1" and "As to claim 4." Further, figure 3 shows a network as part of the system.

As to claim 12, Schimmel teaches that **requesting the operating system enable remote translation passes control of the RTT to the operating system** [whenever a translation in TLB 718 is to be invalidated, CPU 714 or the operating

system must be interrupted in order to execute the invalidation (column 9, lines 65-67); a copy of the translation is also placed in TLB 718. Later, when CPU 714 requires a translation, CPU 714 or an operating system searches TLB 718. If the translation is not found in TLB 718 (i.e., a TLB "miss"), the desired translation can be loaded from the page tables in memory by hardware, software, firmware, or any combination thereof (column 9, lines 22-28); an operating system that controls processor and cache node 410 can map virtual memory, as discussed above in FIGS. 5 and 6, where virtual memory addresses 512 are mapped to physical addresses 516 (column 11, lines 52-54); in step 912, the operating system generates virtual memory address-to-physical memory address translations for the mapped data. For example, a page table, such as page table 610 can be generated for a process that is provided with mapped memory. The page table can be for the exclusive use of one user or can be shared by multiple users (column 11, lines 59-64); the operation can be performed solely by the CPU, solely by an operating system (not shown) that controls the CPU or by a combination of the CPU and the operating system (column 10, lines 39-42); the PTE for page 0x10 can be found by simply looking at index 0x10 in the page table array. The starting address of the array itself is maintained by the operating system in such a way that it is easy to find (column 8, lines 40-44)].

As to claim 13, Schimmel teaches that **passing control of the RTT to the operating system causes the operating system to maintain coherency of the RTT** [a system and method for maintaining consistency between translational look-aside buffers (TLB) and page tables (abstract); one problem that confronts both TLBs and

caches is maintaining consistency of data that is stored in more than one location ...

Thus, the operating system updates the PTE to reflect the new physical location of the data (column 2, lines 12-25); cache consistency can be maintained between data cached in cache 416 -and data stored in main memory 810 by a variety of consistency techniques. For example, cache consistency can be maintained with an optional cache consistency directory 814 in main memory 810. Alternatively, cache consistency can be maintained by a snooping protocol implemented within cache controller 812 which snoops bus 816 for broadcast messages (column 10, lines 15-22); the operation can be performed solely by the CPU, solely by an operating system (not shown) that controls the CPU or by a combination of the CPU and the operating system (column 10, lines 39-42)].

As to claim 14, refer to "As to claim 12."

As to claim 15, refer to "As to claim 13."

As to claim 16, refer to "As to claim 12."

As to claim 17, refer to "As to claim 13."

## 6.

### ***Related Prior Art of Record***

The following list of prior art is considered to be pertinent to applicant's invention, but not relied upon for claim analysis conducted above.

- Scott et al. (US 6,925,547), "Remote Address Translation in a Multiprocessor System."

- Scott, (US 6,922,766), "Remote Translation Mechanism for a Multi-Node System."

***Conclusion***

7. Claims 1-11 are rejected as explained above.

8. **THIS ACTION IS MADE FINAL.** Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the mailing date of this final action.

9. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Sheng-Jen Tsai whose telephone number is 571-272-4244. The examiner can normally be reached on 8:30 - 5:00.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Matthew Kim can be reached on 571-272-4182. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

Sheng-Jen Tsai  
Examiner  
Art Unit 2186

February 25, 2006

  
PIERRE BATAILLE  
PRIMARY EXAMINER  
3/21/06